Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L2	12961	file near4 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:43
L3	0	("220040193672").pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 06:59
L4	2	("20040193672").pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 06:59
L5	1	4 and file near4 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:00
L6	0	4 and query near9 result near9 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:03
L7	1	4 and query near9 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:03
L8	0	4 and query with result with file with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:04
L9	0	4 and query with file with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:04
L10	0	4 and query with (item or file) with result with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:04

L11	1	4 and query with (item or file) with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:05
L12	1	4 and (query or find$3 or request$5 or search$6) with (item or file) with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:39
L13	1	4 and (query or find$3 or request$5 or search$6)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:07
L14	1	4 and (dynamic and static) adj4 (list or directory or folder or container)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:10
L15	1	14 and non adj file	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:22
L16	1	15 and item	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:09
L17	1	4 and scop$3 and criteria and set and (physical adj location) and storage	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:24
L18	1	4 and (dynamic and static and virtual) adj4 (list or directory or folder or container)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:12
L19	31	(dynamic and static and virtual) adj4 (list or directory or folder or container)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:21
L20	2338	(query or find$3 or request$5 or search$6) with (item or file) with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:22

L21	667	(query or find$3 or request$5 or search$6) adj9 (item or file) adj9 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:27
L22	24	21 and non adj file	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:22
L23	0	21 and (scope or (physical adj location) or (data adj storage)) with creteira	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:25
L24	6	21 and (scope or (physical adj location) or (data adj storage)) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:28
L25	294	((query or find$3 or search$6) adj9 (item or file)) adj9 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:27
L26	237	25 and file near4 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:28
L27	198	25 and file adj shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:28
L28	212	25 and file near shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:28
L29	2	28 and (query or find$3 or locate$1) with (scope or (physical adj location) or (data adj storage)) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:32
L30	692	2 and (construct$6 or creat$6 or generat$6 or implement$6 or establish$6 or build$4) adj4 query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:34

L31	255	30 and (query or find$3 or locate$1 or search$5) with shar$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:34
L32	4	31 and (query or find$3 or locate$1) with (scope or (physical adj location) or (data adj storage)) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:35
L33	13611	(construct$6 or creat$6 or generat$6 or implement$6 or establish$6 or build$4) adj4 query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12
L34	747	33 and (query or find$3 or locate$1 or search$5) adj9 shar$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:35
L35	7	34 and ((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:47
L36	19	2 and ((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:41
L37	2	36 and (query or find$3 or search$6) adj9 ((item or file) near6 shar$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:42
L38	2	36 and (query or find$3 or search$6) with ((item or file) near6 shar$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:42
L39	6	36 and (query or find$3 or search$6) with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:42
L40	73	((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:10

L41	2	40 and (query or find$3 or search$6) adj9 ((item or file) near6 shar$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:44
L42	2	40 and (query or find$3 or search$6) with ((item or file) near6 shar$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:42
L43	2	42 and (query or find$3 or search$6) adj20 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:43
L44	2	42 and (query or find$3 or search$6) with shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:43
L45	7586	(query or find$3 or search$6) adj20 shar$5	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:43
L46	963	45 and file near4 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:43
L47	662	45 and file adj3 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:44
L48	378	(query or find$3 or search$6) adj9 ((item or file) near6 shar$5)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:44
L49	273	48 and file adj3 shar$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:44
L50	85	49 and "707"/$.ccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:07

L51	2	49 and ((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:19
L52	2	49 and ((query or find$3 or locate$1 or search$6 or request$6) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:48
L53	52	49 and ((query or find$3 or locate$1 or search$6 or request$6) adj10 (scope or (physical adj location) or (data adj storage))) and criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:48
L54	4	49 and ((query or find$3 or locate$1 or search$6 or request$6) adj10 (scope or (physical adj location) or (data adj storage))) same criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 07:48
L55	27	50 and (p2p or peer-to-peer or gnutella)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:07
L56	8	50 and (p2p or peer-to-peer) and gnutella	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:18
L57	0	56 and ((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12
L58	1	56 and (scope or (physical adj location) or (data adj storage)) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12
L59	301	(p2p or peer-to-peer) and gnutella	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:21
L60	0	59 and ((query or find$3 or locate$1) adj10 (scope or (physical adj location) or (data adj storage))) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12

L61	4	59 and (scope or (physical adj location) or (data adj storage)) with criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12
L62	1	61 and (construct$6 or creat$6 or generat$6 or implement$6 or establish$6 or build$4) adj4 query	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:12
L63	4	("20030225796" or "20020033844").pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:18
L64	1	63 and (query or find$3 or locate$1 or search$6 or request$6) and (scope or (physical adj location) or (data adj storage)) and criteria	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:20
L65	6388	(query or find$3 or locate$1 or search$6 or request$6) and ((scope or (physical adj location) or (data adj storage)) same criteria)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:21
L66	6	65 and (p2p or peer-to-peer) and gnutella	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:24
L67	7	65 and gnutella	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:21
L68	0	65 and gnutella and acl	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:21
L69	2	59 and acl	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	ON	2006/02/09 08:25

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

title: (+p2p or peer-to-peer or gnutella) +sharing +query SEARCH

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **p2p** or **peer to peer** or **gnutella sharing query** Found **13,291** of **169,866**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Peer-to-peer: Making gnutella-like P2P systems scalable

Yatin Chawathe, Sylvia Ratnasamy, Lee Breslau, Nick Lanham, Scott Shenker

August 2003 **Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications**

Publisher: ACM Press

Full text available: pdf(181.52 KB) Additional Information: full citation, abstract, references, citings, index terms

Napster pioneered the idea of peer-to-peer file sharing, and supported it with a centralized file search facility. Subsequent P2P systems like Gnutella adopted decentralized search algorithms. However, Gnutella's notoriously poor scaling led some to propose distributed hash table solutions to the wide-area file search problem. Contrary to that trend, we advocate retaining Gnutella's simplicity while proposing new mechanisms that greatly improve its scalability. Building upon prior research [1, 1 ...

Keywords: Gnutella, distributed hash tables, peer-to-peer

2 Traffic characterization: Characterizing the query behavior in peer-to-peer file sharing systems

Alexander Klemm, Christoph Lindemann, Mary K. Vernon, Oliver P. Waldhorst

October 2004 **Proceedings of the 4th ACM SIGCOMM conference on Internet measurement**

Publisher: ACM Press

Full text available: pdf(525.94 KB) Additional Information: full citation, abstract, references, index terms

This paper characterizes the query behavior of peers in a peer-to-peer (P2P) file sharing system. In contrast to previous work, which provides various aggregate workload statistics, we characterize peer behavior in a form that can be used for constructing representative synthetic workloads for evaluating new P2P system designs. In particular, the analysis exposes heterogeneous behavior that occurs on different days, in different geographical regions (i. e., Asia, Europe, and North America) or ...

Keywords: overlay networks, peer-to-peer, synthetic workloads, workload characterization

3 XML schemas: integration and translation: Intelligent knowledge discovery in peer-to-

peer file sharing
Yugyung Lee, Changgyu Oh, Eun Kyo Park
November 2002 **Proceedings of the eleventh international conference on Information and knowledge management**
Publisher: ACM Press
Full text available: pdf(361.17 KB) Additional Information: full citation, abstract, references, index terms

Emerging peer-to-peer computing provides new possibilities but also challenges for distributed applications. Despite their significant potential, current peer-to-peer networks lack efficient knowledge discovery and management. This paper addresses this deficiency and proposes the Intelligent File Sharing framework, which provides an effective and flexible query for P2P file sharing. The IFS is based on powerful schema and flexible inference, as well as efficiently integrated and extensible retri ...

Keywords: association rules, encoding, hierarchy, peer-to-peer file sharing, reasoning, retrieval, search

4 Interoperability of peer-to-peer file sharing protocols
Siu Man Lui, Sai Ho Kwok
June 2002 **ACM SIGecom Exchanges**, Volume 3 Issue 3
Publisher: ACM Press
Full text available: pdf(42.71 KB) Additional Information: full citation, abstract, references, citings, index terms

Peer-to-Peer (P2P) file sharing software has brought a hot discussion on P2P file sharing among all businesses. Freenet, Gnutella, and Napster are the three most popular P2P file sharing applications. They use three distinct protocols and these protocols come with different characteristics. In this paper, we discuss the protocols of these P2P file sharing applications, in terms of the methodologies used for peer registry, query and content sharing. In order to maximize the benefit of P2P file sh ...

Keywords: Gnutella, Napster, Peer-to-Peer

5 Denial-of-service resilience in peer-to-peer file sharing systems
D. Dumitriu, E. Knightly, A. Kuzmanovic, I. Stoica, W. Zwaenepoel
June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1
Publisher: ACM Press
Full text available: pdf(245.14 KB) Additional Information: full citation, abstract, references, citings, index terms

Peer-to-peer (p2p) file sharing systems are characterized by highly replicated content distributed among nodes with enormous aggregate resources for storage and communication. These properties alone are not sufficient, however, to render p2p networks immune to denial-of-service (DoS) attack. In this paper, we study, by means of analytical modeling and simulation, the resilience of p2p file sharing systems against DoS attacks, in which malicious nodes respond to queries with erroneous responses. ...

Keywords: denial of service, file pollution, network-targeted attacks, peer-to-peer

6 Using locality of reference to improve performance of peer-to-peer applications
Marcelo Werneck Barbosa, Melissa Morgado Costa, Jussara M. Almeida, Virgílio A. F. Almeida
January 2004 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 4th**

international workshop on Software and performance WOSP '04, Volume 29 Issue 1
Publisher: ACM Press
Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, references

Peer-to-peer, or simply P2P, systems have recently emerged as a popular paradigm for building distributed applications. One key aspect of the P2P system design is the mechanism used for content location. A number of different approaches are currently in use. In particular, the location algorithm used in Gnutella, a popular and extensively analyzed P2P file sharing application, is based on flooding of messages in the network, which results in significant processing overhead on the participant nod ...

Keywords: content location algorithms, interest-based communities, peer-to-peer systems, performance analysis

7 Distributed content-based visual information retrieval system on peer-to-peer networks
Irwin King, Cheuk Hang Ng, Ka Cheung Sia
July 2004 **ACM Transactions on Information Systems (TOIS)**, Volume 22 Issue 3
Publisher: ACM Press
Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings, index terms

With the recent advances of distributed computing, the limitation of information retrieval from a centralized image collection can be removed by allowing distributed image data sources to interact with each other for data storage sharing and information retrieval. In this article, we present our design and implementation of DISCOVIR: DIStributed COntent-based Visual Information Retrieval system using the Peer-to-Peer (P2P) Network. We describe the system architecture and detail the interactions ...

Keywords: Peer-to-peer (P2P) network, content-based image retrieval (CBIR), information retrieval, intelligent query routing, peer clustering

8 Content availability, pollution and poisoning in file sharing peer-to-peer networks
Nicolas Christin, Andreas S. Weigend, John Chuang
June 2005 **Proceedings of the 6th ACM conference on Electronic commerce**
Publisher: ACM Press
Full text available: pdf(228.88 KB) Additional Information: full citation, abstract, references, index terms

Copyright holders have been investigating technological solutions to prevent distribution of copyrighted materials in peer-to-peer file sharing networks. A particularly popular technique consists in "poisoning" a specific item (movie, song, or software title) by injecting a massive number of decoys into the peer-to-peer network, to reduce the availability of the targeted item. In addition to poisoning, pollution, that is, the accidental injection of unusable copies of files in the network, also ...

Keywords: content protection, file sharing, peer-to-peer networks

9 Peer-to-peer computing: Foreseer: a novel, locality-aware peer-to-peer system architecture for keyword searches
Hailong Cai, Jun Wang
October 2004 **Proceedings of the 5th ACM/IFIP/USENIX international conference on Middleware**
Publisher: Springer-Verlag New York, Inc.

Full text available: pdf(315.85 KB) Additional Information: full citation, abstract, references

Peer-to-peer (P2P) systems are becoming increasingly popular and complex, serving millions of users today. However, the design of current unstructured P2P systems does not take full advantage of rich locality properties present in P2P system workloads, thus possibly resulting in inefficient searches or poor system scalability. In this paper, we propose a novel locality-aware P2P system architecture called Foreseer, which explicitly exploits <I>geographical</I> locality and <I>t ...

Keywords: Bloom filter, Foreseer, geographical locality, temporal locality

10 A survey of peer-to-peer content distribution technologies



Stephanos Androutsellis-Theotokis, Diomidis Spinellis

December 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 4

Publisher: ACM Press

Full text available: pdf(517.77 KB) Additional Information: full citation, abstract, references, index terms

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance.Content distribution is an important peer-to-peer application ...

Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

11 Protocols: The Eigentrust algorithm for reputation management in P2P networks



Sepandar D. Kamvar, Mario T. Schlosser, Hector Garcia-Molina

May 2003 **Proceedings of the 12th international conference on World Wide Web**

Publisher: ACM Press

Full text available: pdf(202.87 KB) Additional Information: full citation, abstract, references, citings, index terms

Peer-to-peer file-sharing networks are currently receiving much attention as a means of sharing and distributing information. However, as recent experience shows, the anonymous, open nature of these networks offers an almost ideal environment for the spread of self-replicating inauthentic files.We describe an algorithm to decrease the number of downloads of inauthentic files in a peer-to-peer file-sharing network that assigns each peer a unique global trust value, based on the peer's history of ...

Keywords: distributed eigenvector computation, peer-to-peer, reputation

12 Session 5: P2P and streaming: Analyzing peer-to-peer traffic across large networks



Subhabrata Sen, Jia Wang

November 2002 **Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment**

Publisher: ACM Press

Full text available: pdf(1.56 MB) Additional Information: full citation, abstract, references, citings, index terms

The use of peer-to-peer (P2P) applications is growing dramaticaliy, particularly for sharing large video/audio files and software. In this paper, we analyze P2P traffic by measuring flow-level information collected at multiple border routers across a large ISP network, and report our investigation of three popular P2P systems -- FastTrack, Gnutella, and

DirectConnect. We characterize the P2P traffic observed at a single ISP and its impact on the underlying network. We observe very skewed distrib ...

13 Session 4: Limited reputation sharing in P2P systems

Sergio Marti, Hector Garcia-Molina

May 2004 **Proceedings of the 5th ACM conference on Electronic commerce**

Publisher: ACM Press

Full text available: pdf(268.20 KB) Additional Information: full citation, abstract, references, citings, index terms

The increasing popularity of resource exchange through peer-to-peer networks has encouraged the development of ways to support more complex commercial transactions over these networks. Unfortunately, the prospect of higher volume and higher value transactions attracts agents seeking to exploit or weaken the network by propagating bad information and services. This paper presents advantages and disadvantages of resource selection techniques based on peer reputation. We evaluate the effect of limi ...

Keywords: peer-to-peer, reputation, trust

14 Peer to peer systems: A reputation system for peer-to-peer networks

Minaxi Gupta, Paul Judge, Mostafa Ammar

June 2003 **Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video**

Publisher: ACM Press

Full text available: pdf(186.94 KB) Additional Information: full citation, abstract, references, citings, index terms

We investigate the design of a reputation system for decentralized unstructured P2P networks like Gnutella. Having reliable reputation information about peers can form the basis of an incentive system and can guide peers in their decision making (e.g., who to download a file from). The reputation system uses objective criteria to track each peer's contribution in the system and allows peers to store their reputations locally. Reputation are computed using either of the two schemes, *debit-cred* ...

Keywords: *P2P networks, reputation system*

15 Analyzing peer-to-peer traffic across large networks

Subhabrata Sen, Jia Wang

April 2004 **IEEE/ACM Transactions on Networking (TON)**, Volume 12 Issue 2

Publisher: IEEE Press

Full text available: pdf(616.51 KB) Additional Information: full citation, abstract, references, index terms, review

The use of peer-to-peer (P2P) applications is growing dramatically, particularly for sharing large video/audio files and software. In this paper, we analyze P2P traffic by measuring flow-level information collected at multiple border routers across a large ISP network, and report our investigation of three popular P2P systems--FastTrack, Gnutella, and Direct-Connect. We characterize the P2P trafffic observed at a single ISP and its impact on the underlying network. We observe very skewed distrib ...

Keywords: P2P, file sharing, peer-to-peer, traffic characterization, traffic measurement

16 XML schemas: integration and translation: A local search mechanism for peer-to-peer networks

Vana Kalogeraki, Dimitrios Gunopulos, D. Zeinalipour-Yazti
November 2002 **Proceedings of the eleventh international conference on Information and knowledge management**
Publisher: ACM Press

Full text available: pdf(238.97 KB) Additional Information: full citation, abstract, references, citings, index terms

One important problem in peer-to-peer (P2P) networks is searching and retrieving the correct information. However, existing searching mechanisms in pure peer-to-peer networks are inefficient due to the decentralized nature of such networks. We propose two mechanisms for information retrieval in pure peer-to-peer networks. The first, the modified Breadth-First Search (BFS) mechanism, is an extension of the current Gnuttela protocol, allows searching with keywords, and is designed to minimize the ...

Keywords: distributed information retrieval, peer-to-peer netwroks

17 Applications: A cross-layer optimization of gnutella for mobile ad hoc networks
Marco Conti, Enrico Gregori, Giovanni Turi
May 2005 **Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing MobiHoc '05**
Publisher: ACM Press
Full text available: pdf(322.67 KB) Additional Information: full citation, abstract, references, index terms

In recent years, the Internet has witnessed the introduction of many peer-to-peer systems designed to realize large-scale data sharing. These platforms exhibit interesting features like self-configuration, self-healing and complete decentralization, which make them appealing for employment in ad hoc environments as well. However, the impact of ad hoc dynamics on the performance of these protocols, and the different set of constraints which this emerging networking paradigm imposes, haven't been ...

Keywords: ad hoc networks, cross-layer design, peer-to-peer networking

18 Peer to peer networks: Query-flood DoS attacks in gnutella
Neil Daswani, Hector Garcia-Molina
November 2002 **Proceedings of the 9th ACM conference on Computer and communications security**
Publisher: ACM Press

Full text available: pdf(465.15 KB) Additional Information: full citation, abstract, references, citings, index terms

We describe a simple but effective traffic model that can be used to understand the effects of denial-of-service (DoS) attacks based on query floods in Gnutella networks. We run simulations based on the model to analyze how different choices of network topology and application level load balancing policies can minimize the effect of these types of DoS attacks. In addition, we also study how damage caused by query floods is distributed throughout the network, and how application-level policies ca ...

Keywords: denial-of-service, peer-to-peer, security

19 Networks: Search and replication in unstructured peer-to-peer networks
Qin Lv, Pei Cao, Edith Cohen, Kai Li, Scott Shenker
June 2002 **Proceedings of the 16th international conference on Supercomputing**
Publisher: ACM Press

Full text available: pdf(510.93 KB) Additional Information: full citation, abstract, references, citings, index

terms

Decentralized and unstructured peer-to-peer networks such as Gnutella are attractive for certain applications because they require no centralized directories and no precise control over network topology or data placement. However, the flooding-based query algorithm used in Gnutella does not scale; each query generates a large amount of traffic and large systems quickly become overwhelmed by the query-induced load. This paper explores, through simulation, various alternatives to Gnutella's query ...

Keywords: peer-to-peer, replication, search, unstructured

20 IR-6 (information retrieval): digital libraries: The robustness of content-based search in hierarchical peer to peer networks

M. Elena Renda, Jamie Callan

November 2004 **Proceedings of the thirteenth ACM international conference on Information and knowledge management CIKM '04**

Publisher: ACM Press

Full text available: pdf(3.27 MB) Additional Information: full citation, abstract, references, index terms

Hierarchical peer to peer networks with multiple directory services are an important architecture for large-scale file sharing due to their effectiveness and efficiency. Recent research argues that they are also an effective method of providing large-scale content-based federated search of text-based digital libraries. In both cases the directory services are critical resources that are subject to attack or failure, but the latter architecture may be particularly vulnerable bec ...

Keywords: content-based, hierarchical, peer to peer, retrieval, robustness, search

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login

Search: The ACM Digital Library The Guide

title: (+p2p or peer-to-peer or gnutella) +shared +item +quer

THE ACM DIGITAL LIBRARY

Feedback Report a problem Satisfaction survey

Terms used **p2p** or **peer to peer** or **gnutella shared item query** Found **4,904** of **169,866**

Sort results by: relevance

Display results: expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown Relevance scale

1 Content availability, pollution and poisoning in file sharing peer-to-peer networks

Nicolas Christin, Andreas S. Weigend, John Chuang

June 2005 **Proceedings of the 6th ACM conference on Electronic commerce**

Publisher: ACM Press

Full text available: pdf(228.88 KB) Additional Information: full citation, abstract, references, index terms

Copyright holders have been investigating technological solutions to prevent distribution of copyrighted materials in peer-to-peer file sharing networks. A particularly popular technique consists in "poisoning" a specific item (movie, song, or software title) by injecting a massive number of decoys into the peer-to-peer network, to reduce the availability of the targeted item. In addition to poisoning, pollution, that is, the accidental injection of unusable copies of files in the network, also ...

Keywords: content protection, file sharing, peer-to-peer networks

2 Using locality of reference to improve performance of peer-to-peer applications

Marcelo Werneck Barbosa, Melissa Morgado Costa, Jussara M. Almeida, Virgílio A. F. Almeida

January 2004 **ACM SIGSOFT Software Engineering Notes , Proceedings of the 4th international workshop on Software and performance WOSP '04**, Volume 29 Issue 1

Publisher: ACM Press

Full text available: pdf(1.36 MB) Additional Information: full citation, abstract, references

Peer-to-peer, or simply P2P, systems have recently emerged as a popular paradigm for building distributed applications. One key aspect of the P2P system design is the mechanism used for content location. A number of different approaches are currently in use. In particular, the location algorithm used in Gnutella, a popular and extensively analyzed P2P file sharing application, is based on flooding of messages in the network, which results in significant processing overhead on the participant nod ...

Keywords: content location algorithms, interest-based communities, peer-to-peer systems, performance analysis

3 A survey of peer-to-peer content distribution technologies

Stephanos Androutsellis-Theotokis, Diomidis Spinellis

December 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 4

Publisher: ACM Press
Full text available: pdf(517.77 KB) Additional Information: full citation, abstract, references, index terms

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance.Content distribution is an important peer-to-peer application ...

Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

4 Research articles and surveys: Peer-to-peer management of XML data: issues and research challenges
Georgia Koloniari, Evaggelia Pitoura
June 2005 **ACM SIGMOD Record**, Volume 34 Issue 2
Publisher: ACM Press
Full text available: pdf(301.94 KB) Additional Information: full citation, abstract, references, index terms

Peer-to-peer (p2p) systems are attracting increasing attention as an efficient means of sharing data among large, diverse and dynamic sets of users. The widespread use of XML as a standard for representing and exchanging data in the Internet suggests using XML for describing data shared in a p2p system. However, sharing XML data imposes new challenges in p2p systems related to supporting advanced querying beyond simple keyword-based retrieval. In this paper, we focus on data management issues fo ...

5 A case for associative peer to peer overlays
Edith Cohen, Amos Fiat, Haim Kaplan
January 2003 **ACM SIGCOMM Computer Communication Review**, Volume 33 Issue 1
Publisher: ACM Press
Full text available: pdf(749.59 KB) Additional Information: full citation, abstract, references, index terms

The success of a P2P file-sharing network highly depends on the scalability and versatility of its search mechanism. Two particularly desirable search features are scope (ability to find infrequent items) and support for partial-match queries (queries that contain typos or include a subset of keywords). While centralized-index architectures (such as Napster) can support both these features, existing decentralized architectures seem to support at most one: prevailing protocols (such as Gnutella a ...

6 Applications: Fast retrieval of high-dimensional feature vectors in P2P networks using compact peer data summaries
Wolfgang Müller, Andreas Henrich
November 2003 **Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval**
Publisher: ACM Press
Full text available: pdf(378.07 KB) Additional Information: full citation, abstract, references, index terms

The retrieval facilities of most Peer-to-Peer (P2P) systems are limited to queries based on a unique identifier or a small set of keywords. The techniques used for this purpose are hardly applicable for content-based image retrieval (CBIR) in a P2P network. Furthermore, we will argue that the curse of dimensionality and the high communication overhead prevent the adaptation of multidimensional search trees or fast sequential scan techniques for P2P CBIR. In the present paper we will propose two ...

7 Replication strategies in unstructured peer-to-peer networks

Edith Cohen, Scott Shenker
August 2002 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '02**, Volume 32 Issue 4
Publisher: ACM Press

Full text available: pdf(371.85 KB) Additional Information: full citation, abstract, references, citings, index terms

The Peer-to-Peer (P2P) architectures that are most prevalent in today's Internet are decentralized and unstructured. Search is blind in that it is independent of the query and is thus not more effective than probing randomly chosen peers. One technique to improve the effectiveness of blind search is to proactively replicate data. We evaluate and compare different replication strategies and reveal interesting structure: Two very common but very different replication strategies - uniform and propo ...

Keywords: peer-to-peer, random search, replication

8 Session 4: Limited reputation sharing in P2P systems
Sergio Marti, Hector Garcia-Molina
May 2004 **Proceedings of the 5th ACM conference on Electronic commerce**
Publisher: ACM Press

Full text available: pdf(268.20 KB) Additional Information: full citation, abstract, references, citings, index terms

The increasing popularity of resource exchange through peer-to-peer networks has encouraged the development of ways to support more complex commercial transactions over these networks. Unfortunately, the prospect of higher volume and higher value transactions attracts agents seeking to exploit or weaken the network by propagating bad information and services. This paper presents advantages and disadvantages of resource selection techniques based on peer reputation. We evaluate the effect of limi ...

Keywords: peer-to-peer, reputation, trust

9 Applications: A cross-layer optimization of gnutella for mobile ad hoc networks
Marco Conti, Enrico Gregori, Giovanni Turi
May 2005 **Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing MobiHoc '05**
Publisher: ACM Press
Full text available: pdf(322.67 KB) Additional Information: full citation, abstract, references, index terms

In recent years, the Internet has witnessed the introduction of many peer-to-peer systems designed to realize large-scale data sharing. These platforms exhibit interesting features like self-configuration, self-healing and complete decentralization, which make them appealing for employment in ad hoc environments as well. However, the impact of ad hoc dynamics on the performance of these protocols, and the different set of constraints which this emerging networking paradigm imposes, haven't been ...

Keywords: ad hoc networks, cross-layer design, peer-to-peer networking

10 Special topic section on peer to peer data management: P-Grid: a self-organizing structured P2P system
Karl Aberer, Philippe Cudré-Mauroux, Anwitaman Datta, Zoran Despotovic, Manfred Hauswirth, Magdalena Punceva, Roman Schmidt
September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3

Publisher: ACM Press
Full text available: pdf(63.00 KB) Additional Information: full citation, references, citings

11 Research papers: storage, indexing, and system architecture: Guaranteeing correctness and availability in P2P range indices
Prakash Linga, Adina Crainiceanu, Johannes Gehrke, Jayavel Shanmugasudaram
June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**
Publisher: ACM Press
Full text available: pdf(430.28 KB) Additional Information: full citation, abstract, references

New and emerging P2P applications require sophisticated range query capability and also have strict requirements on query correctness, system availability and item availability. While there has been recent work on developing new P2P range indices, none of these indices guarantee correctness and availability. In this paper, we develop new techniques that can provably guarantee the correctness and availability of P2P range indices. We develop our techniques in the context of a general P2P indexing ...

12 Information retrieval session 4: general retrieval issues I: Content-based retrieval in hybrid peer-to-peer networks
Jie Lu, Jamie Callan
November 2003 **Proceedings of the twelfth international conference on Information and knowledge management**
Publisher: ACM Press
Full text available: pdf(262.41 KB) Additional Information: full citation, abstract, references, citings, index terms

Hybrid peer-to-peer architectures use special nodes to provide directory services for regions of the network ("regional directory services"). Hybrid peer-to-peer architectures are a potentially powerful model for developing large-scale networks of complex digital libraries, but peer-to-peer networks have so far tended to use very simple methods of resource selection and document retrieval. In this paper, we study the application of content-based resource selection and document retrieval to hybri ...

Keywords: content-based, hybrid, peer-to-peer, retrieval, search

13 Information access and retrieval (IAR): An architecture for information retrieval over semi-collaborating Peer-to-Peer networks
Iraklis A. Klampanos, Joemon M. Jose
March 2004 **Proceedings of the 2004 ACM symposium on Applied computing**
Publisher: ACM Press
Full text available: pdf(365.46 KB) Additional Information: full citation, abstract, references, index terms

Peer-to-Peer (P2P) networking is aimed at exploiting the potential of widely distributed information pools and its effortless access and retrieval irrespectively of underlying networking protocols, operating systems of devices. However, prohibiting limitations have been identified and perhaps the most important one is the successful location of relevant information sources and the efficient query routing in large, highly distributed P2P networks. In this paper, a novel, cluster-based architectur ...

14 XML and semistructured data querying: XPath lookup queries in P2P networks
Angela Bonifati, Ugo Matrangolo, Alfredo Cuzzocrea, Mayank Jain
November 2004 **Proceedings of the 6th annual ACM international workshop on Web information and data management**

Publisher: ACM Press
Full text available: pdf(263.77 KB) Additional Information: full citation, abstract, references, index terms

We address the problem of querying XML data over a P2P network. In P2P networks, the allowed kinds of queries are usually exact-match queries over file names. We discuss the extensions needed to deal with XML data and XPath queries. A single peer can hold a whole document or a partial/complete fragment of the latter. Each XML fragment/document is identified by a distinct path expression, which is encoded in a distributed hash table. Our framework differs from content-based routing mechanisms, ...

Keywords: P2P networks, XML querying, XPath, distributed XML indexes

15 Papers from Hotnets-II: Should we build Gnutella on a structured overlay?
Miguel Castro, Manuel Costa, Antony Rowstron
January 2004 **ACM SIGCOMM Computer Communication Review**, Volume 34 Issue 1
Publisher: ACM Press
Full text available: pdf(127.00 KB) Additional Information: full citation, abstract, references

There has been much interest in both unstructured and structured overlays recently. Unstructured overlays, like Gnutella, build a random graph and use flooding or random walks on the graph to discover data stored by overlay nodes. Structured overlays assign keys to data items and build a graph that maps each key to a specific node. The structure of the graph enables efficient discovery of data items given their keys but it does not support complex queries.Should we build Gnutella on a structured ...

16 Technical and social components of peer-to-peer computing: Looking up data in P2P systems
Hari Balakrishnan, M. Frans Kaashoek, David Karger, Robert Morris, Ion Stoica
February 2003 **Communications of the ACM**, Volume 46 Issue 2
Publisher: ACM Press
Full text available: pdf(105.04 KB) html(33.79 KB) Additional Information: full citation, abstract, references, citings, index terms

The main challenge in P2P computing is to design and implement a robust and scalable distributed system composed of inexpensive, individually unreliable computers in unrelated administrative domains. The participants in a typical P2P system might include computers at homes, schools, and businesses, and can grow to several million concurrent participants.

17 Special topic section on peer to peer data management: Relational data sharing in peer-based data management systems
Beng Chin Ooi, Yanfeng Shu, Kian-Lee Tan
September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3
Publisher: ACM Press
Full text available: pdf(69.76 KB) Additional Information: full citation, abstract, references, citings

Data sharing in current P2P systems is very much restricted to file-system-like capabilities. In this paper, we present the strategies that we have adopted in our BestPeer project to support more fine-grained data sharing, especially, relational data sharing, in a P2P context. First, we look at some of the issues in designing a peer-based data management system, and discuss some possible solutions to address these issues. Second, we present the design of our first prototype system, PeerDB, and r ...

18 Scaling up the semantic web: Super-peer-based routing and clustering strategies for RDF-based peer-to-peer networks

Wolfgang Nejdl, Martin Wolpers, Wolf Siberski, Christoph Schmitz, Mario Schlosser, Ingo Brunkhorst, Alexander Löser
May 2003 **Proceedings of the 12th international conference on World Wide Web**
Publisher: ACM Press

Full text available: pdf(310.31 KB) Additional Information: full citation, abstract, references, citings, index terms

RDF-based P2P networks have a number of advantages compared with simpler P2P networks such as Napster, Gnutella or with approaches based on distributed indices such as CAN and CHORD. RDF-based P2P networks allow complex and extendable descriptions of resources instead of fixed and limited ones, and they provide complex query facilities against these metadata instead of simple keyword-based searches.In previous papers, we have described the Edutella infrastructure and different kinds of Edutella ...

Keywords: distributed RDF repositories, peer-to-peer, schema-based routing, semantic web

19 Privacy and trust: Usability and privacy: a study of Kazaa P2P file-sharing
Nathaniel S. Good, Aaron Krekelberg
April 2003 **Proceedings of the SIGCHI conference on Human factors in computing systems**
Publisher: ACM Press

Full text available: pdf(444.04 KB) Additional Information: full citation, abstract, references, citings, index terms

P2P file sharing systems such as Gnutella, Freenet, and KaZaA, while primarily intended for sharing multimedia files, frequently allow other types of information to be shared. This raises serious concerns about the extent to which users may unknowingly be sharing private or personal information.In this paper, we report on a cognitive walkthrough and a laboratory user study of the KaZaA file sharing user interface. The majority of the users in our study were unable to tell what files they were sh ...

Keywords: Peer-to-peer networks

20 Special topic section on peer to peer data management: Design issues and challenges for RDF- and schema-based peer-to-peer systems
Wolfgang Nejdl, Wolf Siberski, Michael Sintek
September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3
Publisher: ACM Press

Full text available: pdf(135.94 KB) Additional Information: full citation, abstract, references

Databases have employed a schema-based approach to store and retrieve structured data for decades. For peer-to-peer (P2P) networks, similar approaches are just beginning to emerge. While quite a few database techniques can be re-used in this new context, a P2P data management infrastructure poses additional challenges which have to be solved before schema-based P2P networks become as common as schema-based databases. We will describe some of these challenges and discuss approaches to solve them. ...

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player